



UNITED STATES PATENT AND TRADEMARK OFFICE

1/11
UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/923,157	08/06/2001	Akira Yamamoto	36992.00074 (HAL-ID-169)	1510
30256	7590	03/21/2006	EXAMINER	
SQUIRE, SANDERS & DEMPSEY L.L.P 600 HANSEN WAY PALO ALTO, CA 94304-1043			PATEL, ASHOKKUMAR B	
			ART UNIT	PAPER NUMBER
			2154	

DATE MAILED: 03/21/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/923,157	YAMAMOTO ET AL.	
	Examiner	Art Unit	
	Ashok B. Patel	2154	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 28 December 2005.

2a) This action is FINAL. 2b) This action is non-final.

3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1-28 is/are pending in the application.
4a) Of the above claim(s) 14-21 is/are withdrawn from consideration.
5) Claim(s) _____ is/are allowed.
6) Claim(s) 1-13 and 22-28 is/are rejected.
7) Claim(s) _____ is/are objected to.
8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.

10) The drawing(s) filed on _____ is/are: a) accepted or b) objected to by the Examiner.

 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
a) All b) Some * c) None of:
1. Certified copies of the priority documents have been received.
2. Certified copies of the priority documents have been received in Application No. _____.
3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date 3/4/02
4) Interview Summary (PTO-413)
Paper No(s)/Mail Date. ____.
5) Notice of Informal Patent Application (PTO-152)
6) Other:

DETAILED ACTION

1. Claims 1-28 are subject to examination. Claims 14-21 are cancelled.

Response to Arguments

2. Applicant's arguments with respect to claims 1 and 13 have been considered but are moot in view of the new ground(s) of rejection.

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. Claims 1-7, 10, 12, 22-26 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carini et al. (hereinafter Carini) (US 6, 636, 873 B1). in view of Sicola et al. (hereinafter Sicola) (US 6, 629, 264 B1).

Referring to claim 1,

Carini teaches a system comprising:

A plurality of data centers (col. 5, line 4-20, "As shown therein, the computer system 400 includes one or more mobile device gateway servers 402." And "Functionally, the mobile device gateway server 402 acts as a middle tier or bridge between one or more enterprise databases 404."), including a first data center (Fig. 4, element 404) and a second data center, (Fig. 4, elements 422) each data center comprising:

a storage system;(Fig. 4, element 420 for element 402 and col. 5, line 4-20

“Functionally, the mobile device gateway server 402 acts as a middle tier or bridge between one or more enterprise databases 404.”) and

a host server (col. 5, line 4-20, “As shown therein, the computer system 400 includes one or more mobile device gateway servers 402.” And col. 5, line 4-20, “Functionally, the mobile device gateway server 402 acts as a middle tier or bridge between one or more enterprise databases 404.”, col. 6, line 12-25, “According to an embodiment of the present invention, detailed replication logs and replication statistics may be generated and maintained on the enterprise database 404 for each of the mobile device gateway servers 402 within the system 400. The enterprise database 404 may also include tools for defining and maintaining the mobile device gateway server(s) 402 and usage statistics on usage history of the mobile device 406, 408, 410, 412, 414 that are assigned to each of the mobile device gateway servers 402. Using these tools, a particular mobile device 406, 408, 410, 412, 414 may be assigned to and/or removed from selected mobile device gateway server(s) 402.” Note: Functionality inherently teaches of element 404 including a host server);

a directory server (col. 6, line 12-25, “According to an embodiment of the present invention, detailed replication logs and replication statistics may be generated and maintained on the enterprise database 404 for each of the mobile device gateway servers 402 within the system 400. The enterprise database 404 may also include tools for defining and maintaining the mobile device gateway server(s) 402 and usage statistics on usage history of the mobile device 406, 408, 410, 412, 414 that are

assigned to each of the mobile device gateway servers 402. Using these tools, a particular mobile device 406, 408, 410, 412, 414 may be assigned to and/or removed from selected mobile device gateway server(s) 402." Note: Functionality inherently teaches of both elements 404 and 402 of Fig. 4 including directory servers.);

a plurality of access gateways(Fig 4, element 402 and 404, col. 6, line 43 through col. 7, line 20," One implementation of such a replication process is illustrated in FIG. 5. As shown therein, database schemas including snapshots of the data (e.g., tables) to be downloaded to each the mobile devices MD1, MD2, MD3, MD4, MD5 may be stored on the replication database 420. Each of the snapshots SMD1, SMD2, SMD3, SMD4, SMD5 stored in the replication database 420 may contain the latest data from the enterprise database 404 to be downloaded to each of the mobile devices MD1, MD2, MD3, MD4, MD5 via the concentrator 422, as of the last time the user of such mobile devices MD1, MD2, MD3, MD4, MD5 connected to the replication database 420 through the concentrator 422. When each of the mobile devices MD1, MD2, MD3, MD4, MD5 connects to the concentrator 422, data is synchronized bi-directionally, meaning that the mobile devices MD1, MD2, MD3, MD4, MD5 update the replication database 420 and the replication database 420 updates the mobile devices MD1, MD2, MD3, MD4, MD5 with replication data (e.g., snapshots SMD1, SMD2, SMD3, SMD4, SMD5) obtained from the enterprise database 404. In this manner, the replication database 420 may act as a store-and-forward layer using a symmetric synchronization process. Preferably, a synchronization log is kept of all transactions for each of the mobile devices MD1, MD2, MD3, MD4, MD5 and the enterprise database 404. In the case wherein the connection

between one of the mobile devices MD1, MD2, MD3, MD4, MD5 and the concentrator 422 fails, the mobile device reverts back to its original state prior to the synchronization attempt, to avoid data loss and partial updates. The maximum number of mobile devices MD1, MD2, MD3, MD4, MD5 that may be simultaneously connected to the concentrator 422 may be dependent upon the hardware capacity (e.g., disk space) of the replication database 420. According to one embodiment, the number of mobile devices MD1, MD2, MD3, MD4, MD5 that may be simultaneously connected to the concentrator 422 is 512. Once this maximum number is reached, subsequent requests for connection to the concentrator 422 may be queued. For example, the replication process may be compatible with Oracle Symmetric Replication, as described in Oracle 8 Server Concepts cited above. Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks.” Note: Fig. 4, elements 402 and 404 incorporate inherently “a plurality of access gateways”);

a network interconnecting said plurality of data centers, said directory server and said access gateway; wherein each of said plurality of data centers is accessible from at least one of said plurality of access gateways (Fig 4, elements 416).

Although, Carini teaches at col. 4, line 33-37, “The present invention may also be viewed as a computer network, including an enterprise database; a replication database distinct and remote from the enterprise database, the replication database being

mapped to and maintained as a subset image of the enterprise database through a bi-directional replication process;..." as well as at col. 6, line 12-25," The enterprise database 404 may also include tools for defining and maintaining the mobile device gateway server(s) 402 and usage statistics on usage history of the mobile device 406, 408, 410, 412, 414 that are assigned to each of the mobile device gateway servers 402.", Carini explicitly fails to teach wherein responsive to input received via any of said at least one of plurality of access gateways, any of said plurality of data centers may be configured as a primary (source) of data, and any of said plurality of data canters may be configured as a secondary (target) of data in a copy operation.

Sicola teaches at Fig. 15 and col. 18, line 11-61," The present system 100 includes the capability for simultaneous bi-directional remote data replication which permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the

write over to its partner array controller. For example, assume local host 101 writes to controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site provides very useful functionality." (wherein responsive to input received via any of said at least one of plurality of access gateways, any of said plurality of data centers may be

configured as a primary (source) of data, and any of said plurality of data centers may be configured as a secondary (target) of data in a copy operation.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 18, line 11-61), "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503.", and for two, which is useful to Carini when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks."

Referring to claim 2,

Keeping in mind the teachings of Carini as stated above including information about said first data center and said second data center is fetched from said directory server (col. 6, line 12-25), Carini fails to explicitly teach wherein, responsive to said input received via any of plurality of access gateways and thereupon, said first data center maybe configured as a primary (source) of data, and said second data center may be configured as a secondary (target) of data in a copy operation.

Sicola teaches at Fig. 15 and col. 18, line 11-61, "The present system 100 includes the capability for simultaneous bi-directional remote data replication which

permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the write over to its partner array controller. For example, assume local host 101 writes to controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same

volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site provides very useful functionality." (wherein, responsive to said input received via any of plurality of access gateways and thereupon, said first data center maybe configured as a primary (source) of data, and said second data center may be configured as a secondary (target) of data in a copy operation.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that , for one, it offers, as stated by Sicola (col. 18, line 11-61) , "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503.", and for two, which is useful to Carini when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-

based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks."

Referring to claim 3,

Keeping in mind the teachings of Carini as stated above, Carini fails to teach the system of claim 2, wherein, responsive to a second input received via any of said at least one of plurality of access gateways, said first data center may be reconfigured as a secondary (target) of data, and said second data center may be configured as a primary (source) of data in a second copy operation.

Sicola teaches at Fig. 15 and col. 18, line 11-61, "The present system 100 includes the capability for simultaneous bi-directional remote data replication which permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the write over to its partner array controller. For example, assume local host 101 writes to

controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site provides very useful functionality." (wherein, responsive to a second input received via any of said at least one of plurality of access gateways, said first data center may be reconfigured as a secondary (target) of data, and said second data center may be configured as a primary (source) of data in a second copy operation.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 18, line 11-61), "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503.", and for two, which is useful to Carini when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks."

Referring to claim 4,

Carini teaches the system of claim 3, wherein copy operations are synchronous said first data center updating contents of storage from contents of a cache memory prior to being reconfigured to as a secondary (target) in said second copy operation. (co. 7, line 17-20, col. 6, line 61-63, Cache is inherent.).

Referring to claim 5,

Carini teaches the system of claim 2, wherein said information fetched from said directory server comprises proximity information for a source (col. 6, line 12-25) of said input received via said at least one of a plurality of access gateways, (col. 6, line 12-25, Fig 4, element 402 and 404, col. 6, line 43 through col. 7, line 20,").

Carini fails to teach wherein said first data center is configured as a primary (source) of data, and said second data center is configured as a secondary (target) of data in said copy operations.

Sicola teaches at Fig. 15 and col. 18, line 11-61, "The present system 100 includes the capability for simultaneous bi-directional remote data replication which permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the write over to its partner array controller. For example, assume local host 101 writes to controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs

B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site provides very useful functionality." (wherein said first data center is configured as a primary (source) of data, and said second data center is configured as a secondary (target) of data in said copy operations.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 18, line 11-61), "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503." And for two, which is useful to Carini

when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks."

Referring to claim 6,

Keeping in mind the teachings of Carini having multiple data centers (a third data center) as stated above, Carini fails to teach the system of claim 2, wherein said plurality of data centers further comprises a third data center, said third data center being configured as another secondary (target) of data in a copy operation.

Sicola teaches at Fig. 15 and col. 18, line 11-61 and at col. 7, line 7-11, "multiple data center being configured as another secondary (target) of data in a copy operation." (third data center being configured as another secondary (target) of data in a copy operation.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 18, line 11-61), "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503." And for two, which is useful to Carini when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in

the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks.”

Referring to claim 7,

Carini teaches the system of claim 1, further comprising a network interface that provides connection between at least one of plurality of access gateways and a use terminal. (col. 6, line 42- 53).

Referring to claim 10,

Carini teaches the system of claim 1, wherein a storage volume from said first data center and a storage volume from said second data center comprise a copy volume group. (col. 6 , line 42- 63).

Referring to claim 12,

Carini teaches the system of claim 1, said host server further comprising a copy volume group interface process, a read request issue process, and a write request issue process. (col. 6 , line 42- 63, Abstract, col. 7, line 17-20)

Referring to claim 22,

Carini teaches a system comprising:

A plurality of data centers (col. 5, line 4-20, “As shown therein, the computer system 400 includes one or more mobile device gateway servers 402.” And “Functionally, the mobile device gateway server 402 acts as a middle tier or bridge between one or more enterprise databases 404.”), including a first data center (Fig. 4,

element 404) and a second data center, (Fig. 4, elements 422) each data center comprising:

a storage system;(Fig. 4, element 420 for element 402 and col. 5, line 4-20
“Functionally, the mobile device gateway server 402 acts as a middle tier or bridge
between one or more enterprise databases 404.”) and

a host server (col. 5, line 4-20, “As shown therein, the computer system 400
includes one or more mobile device gateway servers 402.” And col. 5, line 4-20,
“Functionally, the mobile device gateway server 402 acts as a middle tier or bridge
between one or more enterprise databases 404.”, col. 6, line 12-25, “According to an
embodiment of the present invention, detailed replication logs and replication statistics
may be generated and maintained on the enterprise database 404 for each of the
mobile device gateway servers 402 within the system 400. The enterprise database
404 may also include tools for defining and maintaining the mobile device gateway
server(s) 402 and usage statistics on usage history of the mobile device 406, 408, 410,
412, 414 that are assigned to each of the mobile device gateway servers 402. Using
these tools, a particular mobile device 406, 408, 410, 412, 414 may be assigned to
and/or removed from selected mobile device gateway server(s) 402.” Note:
Functionality inherently teaches of element 404 including a host server);

a directory server (col. 6, line 12-25, “According to an embodiment of the present
invention, detailed replication logs and replication statistics may be generated and
maintained on the enterprise database 404 for each of the mobile device gateway
servers 402 within the system 400. The enterprise database 404 may also include tools

for defining and maintaining the mobile device gateway server(s) 402 and usage statistics on usage history of the mobile device 406, 408, 410, 412, 414 that are assigned to each of the mobile device gateway servers 402. Using these tools, a particular mobile device 406, 408, 410, 412, 414 may be assigned to and/or removed from selected mobile device gateway server(s) 402." Note: Functionality inherently teaches of both elements 404 and 402 of Fig. 4 including directory servers.);

a plurality of access gateways(Fig 4, element 402 and 404, col. 6, line 43 through col. 7, line 20," One implementation of such a replication process is illustrated in FIG. 5. As shown therein, database schemas including snapshots of the data (e.g., tables) to be downloaded to each the mobile devices MD1, MD2, MD3, MD4, MD5 may be stored on the replication database 420. Each of the snapshots SMD1, SMD2, SMD3, SMD4, SMD5 stored in the replication database 420 may contain the latest data from the enterprise database 404 to be downloaded to each of the mobile devices MD1, MD2, MD3, MD4, MD5 via the concentrator 422, as of the last time the user of such mobile devices MD1, MD2, MD3, MD4, MD5 connected to the replication database 420 through the concentrator 422. When each of the mobile devices MD1, MD2, MD3, MD4, MD5 connects to the concentrator 422, data is synchronized bi-directionally, meaning that the mobile devices MD1, MD2, MD3, MD4, MD5 update the replication database 420 and the replication database 420 updates the mobile devices MD1, MD2, MD3, MD4, MD5 with replication data (e.g., snapshots SMD1, SMD2, SMD3, SMD4, SMD5) obtained from the enterprise database 404. In this manner, the replication database 420 may act as a store-and-forward layer using a symmetric synchronization process. Preferably, a

synchronization log is kept of all transactions for each of the mobile devices MD1, MD2, MD3, MD4, MD5 and the enterprise database 404. In the case wherein the connection between one of the mobile devices MD1, MD2, MD3, MD4, MD5 and the concentrator 422 fails, the mobile device reverts back to its original state prior to the synchronization attempt, to avoid data loss and partial updates. The maximum number of mobile devices MD1, MD2, MD3, MD4, MD5 that may be simultaneously connected to the concentrator 422 may be dependent upon the hardware capacity (e.g., disk space) of the replication database 420. According to one embodiment, the number of mobile devices MD1, MD2, MD3, MD4, MD5 that may be simultaneously connected to the concentrator 422 is 512. Once this maximum number is reached, subsequent requests for connection to the concentrator 422 may be queued. For example, the replication process may be compatible with Oracle Symmetric Replication, as described in Oracle 8 Server Concepts cited above. Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks.” Note: Fig. 4, elements 402 and 404 incorporate inherently “a plurality of access gateways”);

a network interconnecting said plurality of data centers, said directory server and said access gateway; wherein each of said plurality of data centers is accessible from at least one of said plurality of access gateways (Fig 4, elements 416).

Although, Carini teaches at col. 4, line 33-37, "The present invention may also be viewed as a computer network, including an enterprise database; a replication database distinct and remote from the enterprise database, the replication database being mapped to and maintained as a subset image of the enterprise database through a bi-directional replication process;..." as well as at col. 6, line 12-25," The enterprise database 404 may also include tools for defining and maintaining the mobile device gateway server(s) 402 and usage statistics on usage history of the mobile device 406, 408, 410, 412, 414 that are assigned to each of the mobile device gateway servers 402.", Carini explicitly fails to teach wherein responsive to input received via any of plurality of access gateways, any of the plurality of data centers is operable to initiate a bi-directional data copying with any other of the plurality of data centers.

Sicola teaches at Fig. 15 and col. 18, line 11-61," The present system 100 includes the capability for simultaneous bi-directional remote data replication which permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home

and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the write over to its partner array controller. For example, assume local host 101 writes to controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site provides very useful functionality." (wherein responsive to input received via any of

plurality of access gateways, any of the plurality of data centers is operable to initiate a bi-directional data copying with any other of the plurality of data centers.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 18, line 11-61), "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503.", and for two, which is useful to Carini when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks."

Referring to claim 23 and 24,

Carini teaches the directory server is operable to furnish information on the first data center and the second data center (col. 6, line 12-25). Carini teaches wherein the furnished information is indicative of a proximity of the one of the plurality of access gateways associated with the received input and at least one data center and wherein the first data center is selected based on the furnished information (col. 6, line 12-25, Fig 4, element 402 and 404, col. 6, line 43 through col. 7, line 20,").

Referring to claim 25,

Keeping in mind the teachings of Carini as stated above, Carini fails to teach wherein each of the plurality of data centers is operable to act as a primary (source) or a secondary (target) in the data copying operation.

Sicola teaches at Fig. 15 and col. 18, line 11-61, "The present system 100 includes the capability for simultaneous bi-directional remote data replication which permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the write over to its partner array controller. For example, assume local host 101 writes to controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs

B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site provides very useful functionality." (wherein each of the plurality of data centers is operable to act as a primary (source) or a secondary (target) in the data copying operation.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 18, line 11-61), "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503." And for two, which is useful to Carini

when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks."

Referring to claim 26,

Carini teaches proximity information received from the directory server in response to the input. (col. 6, line 12-25, Fig 4, element 402 and 404, col. 6, line 43 through col. 7, line 20).

Carini fails to teach wherein each of the plurality of data centers is configurable as the primary (source) or the secondary (target) in response to the input.

Sicola teaches at Fig. 15 and col. 18, line 11-61, "The present system 100 includes the capability for simultaneous bi-directional remote data replication which permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home

and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the write over to its partner array controller. For example, assume local host 101 writes to controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site

provides very useful functionality." (wherein each of the plurality of data centers is configurable as the primary (source) or the secondary (target) in response to the input.) Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 18, line 11-61), "remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503." And for two, which is useful to Carini when Carini offers, as stated by Carini (col. 7, line 17-20), "Alternatively, the replication process may be configured to be compatible with most any replication protocol known in the art. For example, instead of store-and-forward based synchronization, a message-based synchronization may be implemented, with the mobile device gateway server 402 performing conflict resolution tasks."

5. Claims 8, 9, 11, 27 and 28 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carini et al. (hereinafter Carini) (US 6, 636, 873 B1) and Sicola et al. (hereinafter Sicola) (US 6, 629, 264 B1) as applied to claim 1 above, and further in view of Hale et al. (hereinafter Hale) (US 2002/0184516 A1)

Referring to claim 8,

Carini teaches the storage system (Fig. 4, element 420 for element 402 and col. 5, line 4-20 "Functionally, the mobile device gateway server 402 acts as a middle tier or bridge between one or more enterprise databases 404."), however, Carini and Sicola

both fail to teach system of claim 1, wherein information associated with a virtual volume is stored in a plurality of real volumes in said storage system.

Hale teaches "a virtual volume is stored in a plurality of real volumes in said storage system. (Fig. 1, page 2, para.[0018], "The virtual volume 102 comprises virtual files and/or virtual directories and one or more real volumes 104A-104B.")

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to implement the teachings of Hale to the combined system of Carini and Sicola since the virtual name can be used to represent the real object, file or data, allowing the user know the virtual name and the mapping of the real files, data or objects be kept transparent to the user (page 2, para.[0018] of Hale).

This would have also been obvious because it allows receiving a request to access a virtual volume with a virtual name; mapping the virtual name to the real object; and providing the real object. The method and system uses virtual objects which map to real objects in a computer system. This mapping is transparent to the subject. In this manner, security policies can be enforced over objects stored in file systems without regard to the policies of the file systems. The system can also be used as a gateway to remote file systems built on top of existing file systems. These advantages provide more flexibility in controlling a subject's access to real objects."(Abstract of Hale).

Referring to claim 9,

Carini and Sicola both fail to teach the system of claim 8, wherein a correspondence between said virtual volume and said plurality of real volumes in said storage system is stored in said directory server.

Hale teaches "the system of claim 8, wherein a correspondence between said virtual volume and said plurality of real volumes in said storage system is stored in said directory server. (Figs. 1, 4 and 5 page 2, para.[0018], "The virtual volume 102 comprises virtual files and/or virtual directories and one or more real volumes 104A-104B."(the system of claim 8, wherein a correspondence between said virtual volume and said plurality of real volumes in said storage system is stored in said directory server.)

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to implement the teachings of Hale to the combined system of Carini and Sicola since the virtual name can be used to represent the real object, file or data, allowing the user know the virtual name and the mapping of the real files, data or objects be kept transparent to the user (page 2, para.[0018] of Hale).

This would have also been obvious because it allows receiving a request to access a virtual volume with a virtual name; mapping the virtual name to the real object; and providing the real object. The method and system uses virtual objects which map to real objects in a computer system. This mapping is transparent to the subject. In this manner, security policies can be enforced over objects stored in file systems without regard to the policies of the file systems. The system can also be used as a gateway to remote file systems built on top of existing file systems. These advantages

provide more flexibility in controlling a subject's access to real objects."(Abstract of Hale).

Referring to claim 11,

Carini and Sicola both fail to teach the system of claim 1, said directory server further comprising a log in process and a virtual volume information.

Hale teaches "the system of claim 1, said directory server further comprising a log in process and a virtual volume information (Fig. 2, element 202, para.[0018], [0021]-[0023]).

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to implement the teachings of Hale to the combined system of Carini and Sicola since the virtual name can be used to represent the real object, file or data, allowing the user know the virtual name and the mapping of the real files, data or objects be kept transparent to the user (page 2, para.[0018] of Hale).

This would have also been obvious because it allows receiving a request to access a virtual volume with a virtual name; mapping the virtual name to the real object; and providing the real object. The method and system uses virtual objects which map to real objects in a computer system. This mapping is transparent to the subject. In this manner, security policies can be enforced over objects stored in file systems without regard to the policies of the file systems. The system can also be used as a gateway to remote file systems built on top of existing file systems. These advantages provide more flexibility in controlling a subject's access to real objects."(Abstract of Hale).

Referring to claims 27 and 28,

Keeping in mind the teachings of Carini wherein Carini teaches the system of claim 22, wherein data stored in the storage system of each of the plurality of data centers (col. 5, line 4-20, "As shown therein, the computer system 400 includes one or more mobile device gateway servers 402." And "Functionally, the mobile device gateway server 402 acts as a middle tier or bridge between one or more enterprise databases 404."), however, Carini and Sicola both fail to teach data centers is organized into at least one virtual volume, and stores information associating the at least one virtual volume with a corresponding physical volume.

Hale teaches data center is organized into at least one virtual volume, and stores information associating the at least one virtual volume with a corresponding physical volume. (Fig. 1, page 2, para.[0018], » The virtual volume 102 comprises virtual files and/or virtual directories and one or more real volumes 104A-104B.")

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to implement the teachings of Hale to the combined system of Carini and Sicola since the virtual name can be used to represent the real object, file or data, allowing the user know the virtual name and the mapping of the real files, data or objects be kept transparent to the user (page 2, para.[0018] of Hale).

This would have also been obvious because it allows receiving a request to access a virtual volume with a virtual name; mapping the virtual name to the real object; and providing the real object. The method and system uses virtual objects which map to real objects in a computer system. This mapping is transparent to the subject. In

this manner, security policies can be enforced over objects stored in file systems without regard to the policies of the file systems. The system can also be used as a gateway to remote file systems built on top of existing file systems. These advantages provide more flexibility in controlling a subject's access to real objects."(Abstract of Hale).

6. Claim 13 is rejected under 35 U.S.C. 103(a) as being unpatentable over Hale et al. (hereinafter Hale) (US 2002/0184516 A1) in view of Sicola et al. (hereinafter Sicola) (US 6, 629, 264 B1).

Referring to claim 13,

Hale teaches a method, comprising: receiving a virtual volume name and network interface ID for a user; finding a virtual volume corresponding to said virtual volume name and network interface ID (para. [0021], Fig. 2); selecting a real volume information corresponding to a data center to which said user is logged into (para.[0023]);

Hale fails to teach determining whether said data center is primary; if said data center does not contain primary volume, issuing a request to change a volume within said data center to be primary; and returning a real volume information for said volume within said data center set to primary.

Sicola teaches at Fig. 15 and col. 18, line 11-61," The present system 100 includes the capability for simultaneous bi-directional remote data replication which permits the system to operate in an 'extended cluster' mode, as if each of the remote storage arrays were local relative to the respective remote host. As discussed in the

previous section, FIG. 15 shows array controller pair A1/A2 as the local, 'home' node 1510, and array controller pair B1/B2 as the remote, 'alternate' node 1511. The use of the home/alternate node mechanism adds to the fault tolerant characteristics of the cluster because now if a pair of arrays becomes inoperative for whatever reason, the host on the other side of dual fabric 1503 can still maintain access to its storage. In bidirectional copy mode, the home/alternate node concept is extended from that described in the previous section, in that, now each host computer has its own home and alternate node. If R/W (read and write) access is enabled on an array controller, then the controller will accept writes from any host, and then synchronously send the write over to its partner array controller. For example, assume local host 101 writes to controller A1/A2 after taking a lock out on the appropriate LUN(s). When controller (pair) A1/A2 receives the write request, A1/A2 proceeds to copy the write data to remote controller pair B1/B2. After A1/A2 has received 'success' status from B1/B2's reception of the write data, A1/A2 sends completion status to host 101. At some point, when host 102 writes to B1/B2, a similar, but reverse action occurs with respect to controller pairs B1/B2 and A1/A2. Locking of local LUNs by a host is performed in bidirectional copy mode because of the particular implementation of home/alternate node in this mode of operation, as explained above. The present system achieves bidirectional copy operation by establishing a 'peer' relationship between array controller pairs on opposite sides of link 1503. Both array controller pairs A1/A2 and B1/B2 'see' the same volume name on opposite sides of the link, and both operate as if each one were 'master' with respect to a particular volume. In the present exemplary embodiment of

system 100, remote copy set LUNs have both 'peer controller pairs' A1/B1 and A2/B2 as simultaneous initiator and target, so that writes to either side will be synchronously mirrored to an array controller's partner on the other side of dual fabric 1503. Furthermore, a LUN on one side of fabric 1503 can be set to read-only access, while its remote partner LUN can be set to R/W access. The combination of bi-directional copy with the capability of R/W access at one site and read-only access at the other site provides very useful functionality." (determining whether said data center is primary; if said data center does not contain primary volume, issuing a request to change a volume within said data center to be primary; and returning a real volume information for said volume within said data center set to primary.)"

Therefore it would have been obvious to one of ordinary skill in this art at the time the invention was made to apply the bi-directional synchronization technique of Sicola such that, for one, it offers, as stated by Sicola (col. 3, line 49-52), "maintains consistency between volumes, the system of the present invention provides a mechanism of associating a set of volumes to synchronize the logging to the set of volumes so that when the log is consistent when it is "played back" to the remote site."

This would have been also obvious since, as stated by Sicola at col. 3, line 19-27, "The present system includes an additional novel aspect of grouping logical units, into 'association sets', for logging and failover purposes. The concept of association sets allows the present system provides for proper ordering of I/O operations during logging across multiple volumes. In addition, association sets are employed by the present invention to provide failure consistency by causing the group of logical

units/volumes to all fail at the same time, ensuring a point in time consistency on the remote site."

Conclusion

Examiner's note: Examiner has cited particular columns and line numbers in the references as applied to the claims above for the convenience of the applicant. Although the specified citations are representative of the teachings of the art and are applied to the specific limitations within the individual claim, other passages and figures may apply as well. It is respectfully requested from the applicant in preparing responses, to fully consider the references in entirety as potentially teaching all or part of the claimed invention, as well as the context of the passage as taught by the prior art or disclosed by the Examiner.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Ashok B. Patel whose telephone number is (571) 272-3972. The examiner can normally be reached on 8:00am-5:00pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John A. Follansbee can be reached on (571) 272-3964. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only.

For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Abp


JOHN FOLLANSBEE
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100